REFORM

Newton

AI and the productivity revolution in government

February 2025

Sean Eke

Joe Hill

ABOUT REFORM

Reform is established as the leading Westminster think tank for public service reform. We believe that the State has a fundamental role to play in enabling individuals, families and communities to thrive. But our vision is one in which the State delivers only the services that it is best placed to deliver, within sound public finances, and where both decision-making and delivery is devolved to the most appropriate level. We are committed to driving systemic change that will deliver better outcomes for all.

We are determinedly independent and strictly non-party in our approach. This is reflected in our cross-party Advisory Board and our events programme which seeks to convene likeminded reformers from across the political spectrum.

Reform is a registered charity, the *Reform* Research Trust, charity no. 1103739.

ABOUT REIMAGINING THE STATE

After a decade of disruption, the country faces a moment of national reflection. For too long, Britain has been papering over the cracks in an outdated social and economic model, but while this may bring temporary respite, it doesn't fix the foundations. In 1942 Beveridge stated: "a revolutionary moment in the world's history is a time for revolutions, not for patching." 80 years on, and in the wake of a devastating national crisis, that statement once again rings true. Now is the time to fix Britain's foundations.

Reform's new programme, *Reimagining the State*, will put forward a bold new vision for the role and shape of the State. One that can create the conditions for strong, confident communities, dynamic, innovative markets, and transformative, sustainable public services.

Reimagining Whitehall is one of the major work streams within this programme. This paper is part of our Reimagining Whitehall series.

ACKNOWLEDGEMENTS

Reform would like to thank Newton for kindly supporting this paper and the series of senior-level roundtables which informed it. We would also like to thank all attendees at these four roundtables, in particular the four speakers:

- 'Productivity through potential: the role of AI' with Dr Laura Gilbert, Head of AI for Government, Ellison Institute of Technology and former Chief Analyst and Director of Data Science, 10 Downing Street.
- 'Sweating the assets: AI-enabled planning' with Lieutenant General Sir Tom Copinger-Symes, Deputy Commander of Strategic Command, Ministry of Defence.
- 'AI on the frontline: better decision-making' with Richard Corbridge, Chief Information Officer, SERGO and former Director General, Chief Digital and Information Officer, Department for Work and Pensions.
- 'Breaking the backlogs: managing demand through AI' with Ben Henshall, Deputy Director, Data Science and AI, 10 Downing Street.

Table of contents

1. OVERVIEW…… 6

2. AN INNOVATION MINDSET……...8

2.1 Fitting AI into government……8

2.2 Risk tolerance……...10

3. THE SKILLS TO SUCCEED……...11

3.1 Skills……11

3.2 Leadership……12

3.3 Frontline users…… 12

3.4 Trust and ethics……13

4. STREAMLINING PROCESSES…… 14

4.1 Governing…….. 14

4.2 Funding…….. 15

4.3 Partnering with the private sector……15

5. FIXING THE PLUMBING……17

5.1 Iterative improvements…….. 17

6. CONCLUSION…….. 19

BIBLIOGRAPHY……... 20

EIGHT IDEAS FOR USING AI TO DRIVE THE PRODUCTIVITY REVOLUTION IN GOVERNMENT

Idea 1: Require all departments to report on their plans to adopt AI as part of each department's Outcome Delivery Plan, with priority areas and resources being put into delivering the transformation

Idea 2: Create a cross-government competition to identify the best new ideas for adopting AI in Whitehall, building on work already done by 10 Downing Street's Data Science team (10DS).

Idea 3: Introduce a standardised end-of-pilot evaluation process which determines the lessons learned and whether the pilot should advance to wider rollout.

Idea 4: Increase pay allowances when recruiting into AI-related and AI-adjacent positions.

Idea 5: Produce a standardised AI Essentials training course for Grade 6 and 7 civil servants in non-specialist roles covering basic skills and ethical considerations of AI.

Idea 6: Set aside funding for scaling up successful AI pilots rather than investing all funding up-front in separate projects. Initial funding for pilots should be conditional on the existence of an evaluation plan, clear success measures and requirements for scaling-up.

Idea 7: Streamline the procurement process to make it easier for government to access off-the-shelf products and for SMEs to offer their services to government.

Idea 8: Treat legacy system updates and data quality and access improvements as iterative, agile processes to be done alongside expanding the use of AI in frontline services.

1. Overview: Whitehall needs to reform itself to embrace the power of AI

Artificial Intelligence (AI) has begun to dramatically alter how parts of the UK public sector operate. Social workers in Kingston upon Thames are using AI to halve the time it takes to complete case notes and assessments.[1] Customer service workers at Derby City Council are using an AI assistant capable of handling more than 40 per cent of inbound calls.[2] Doctors at Addenbrooke Hospital are using AI to plan for radiotherapy treatments two and half times faster than if they were working alone.[3] Officers in the British Army are using AI to automate elements of the military operational planning process.[4]

As recognised in the *AI Opportunities Action Plan*, adopting AI more widely across the public sector could lead to more efficient and more effective public services.[5] Examples include:

- **Planning:** Analysing and predicting demand for public services, spotting trends which humans cannot and allowing for more preventative measures and efficient resource allocation.
- **Streamlining:** Automating certain basic tasks and assisting humans with more complex tasks by summarising documents and recommending options.
- **Evaluating:** Running complex policy evaluations, allowing for iterative improvements.

And yet despite the potential for productivity benefits worth, by one estimate, £200 billion over the next five years, AI's deployment throughout the public sector remains patchy and too reliant on individual public servants willing to champion the technology to their leaders and budget-holders.[6]

The problem is not primarily a lack of funding. Large Language Models are prohibitively expensive to develop but cheap to use off-the-shelf, and there is a plethora of low-hanging fruit for the government to pick in areas which could use even cheaper, more tried-and-tested machine learning software. Nor is the problem primarily a lack of regulatory and/or legal clarity. A large corpus of AI guidance now exists for most areas of the public sector.

Instead, the problem is mainly cultural. Government departments and public bodies are not set up to maximise the benefits of innovations like AI. There are few incentives to shift towards AI-enabled delivery models, an aversion towards critical evaluation, limited access to specialist skills, and bureaucratic processes fail to quickly translate ideas into new products and tools. If AI pilots succeed they are not consistently scaled up; if AI pilots fail they are not adequately learned from.

[1] Local Government Association and NHS England, 'Artificial Intelligence Use Cases', Web Page, n.d.
[2] Derby City Council, 'Next Steps for AI Technology to Protect Services for Derby Residents', Web Page, 18 December 2023.
[3] NHS Cambridge University Hospitals, 'AI Cuts Waiting Times for Cancer Patients in NHS First', Web Page, 27 June 2023.
[4] Ministry of Defence, *The Defence AI Playbook*, 2024.
[5] Matt Clifford, *AI Opportunities Action Plan* (Department for Science, Innovation and Technology, 2025).
[6] Alexander Iosad, David Railton, and Tom Westgarth, *Governing in the Age of AI: A New Model to Transform the State* (Tony Blair Institute for Global Change, 2024).

Like other fundamental technological transformations, government cannot expect AI to simply bolt onto the existing machinery of the State. Instead, before the technical transformation can be achieved, there needs to be a cultural transformation.

Across the public sector attitudes towards AI must shift. Leaders need to take charge of this shift, ensuring that innovations are widely adopted, and look afresh at what the functions of the State are and how AI can best enable them. Cross-functional teams collaborating across policy, operations and technology will be essential in embedding this shift and getting real-world value out of AI.

This paper is informed by a series of roundtable discussions with senior figures at the leading edge of AI adoption within government. It discusses the experiences and reflections from these early AI projects and presents a set of policy ideas for enabling a cultural transformation.

These discussions identified four cross-cutting opportunities:

- **Mindset:** fostering an innovative mindset geared towards AI's potential.
- **People:** identifying the right skills, leadership, practices and approach towards trust and ethics.
- **Processes:** streamlining governance, operational guidance, funding and procurement.
- **Plumbing:** fixing the underlying problems of legacy systems and data quality and access.

2. An innovation mindset

There are few things harder for an established organisation — particularly one the size of government — to do than change its culture. By its very nature, culture is intangible, hard to define, and hard to shape. And yet if government is to realise AI's benefits, then changing attitudes towards AI is paramount.

2.1 Fitting AI into government

First thought, not afterthought

If Whitehall is to realise AI's benefits, AI cannot be treated as an afterthought. Attendees described how in the civil service AI does not feature in the early stages of policymaking. The normal processes are followed, and only during the much later stages of implementation is any thought given to what role AI could play.

Devising policy in this way, with insufficient attention given to radically different delivery models enabled by AI, restricts AI to a small number of use cases. It also increases the risk that these use cases fail, because they are unlikely to have been sufficiently thought through. Instead, AI must be considered throughout the policy development cycle to ensure that all the right opportunities are identified and properly scoped.

Science experiment

Similarly, attendees agreed that "AI will fail if it remains only a science experiment". Several commented that in some cases AI was being used in ways which were largely disconnected from the parts of government where it could best help. Instead it can be pigeonholed as something for the digital, data and technology experts to experiment with. And projects are often stuck in a perpetual pilot stage, failing to ever achieve the critical momentum to scale up into day-to-day use and therefore failing, as argued in the *AI Opportunities Action Plan*, to have a meaningful impact on productivity, effectiveness and citizen experience.[7]

Experimenting with AI beyond the normal day-to-day business of government — for example through a 'sandbox' system where teams would develop projects beyond Whitehall's traditional confines — can be a highly effective way of encouraging innovation. However, such experimentation needs to be done with a clear plan to implement prototypes in the wider civil service if they are successful. It cannot be an end in itself, one which costs money and fails to deliver results but still, as one attendee characterised it, "lets policymakers tick a box which says 'I tried to use AI'".

No silver bullets

But attendees were also worried about policymakers giving too much credit to AI too soon: viewing it as the "holy grail" or as a "silver bullet" for all problems. The civil servants we spoke to were very optimistic about AI's potential impact on public services. But they recognised that AI was not easily deployable across all public services in the same way, could not be integrated into all delivery systems, and would not be able to solve every problem. They thought attempting to "shoehorn" AI into policy areas which are not yet ready risked undermining

[7] Clifford, *AI Opportunities Action Plan*.

confidence in the technology, and ultimately preventing it being used in the areas it could do the most good.

In particular, leaders who appear to place excessive faith in AI's suitability for every use case "set themselves up for failure" when some turn out to be unfeasible. This could dent trust in and enthusiasm for AI, and thus consequently slow down public sector adoption. Government needs a nuanced approach, with an awareness of AI's merits balanced against a recognition that it cannot solve everything.

Humans with AI

Part of this nuanced approach involves recognising how public servants would need to work with AI in practice. Multiple use cases exist — cancer screening, large-scale wargame simulations and asylum assessment streamlining — where AI's relevance is clear. But appreciating the importance of human intelligence — grounded in experience and judgement, and better able to consider outliers — is vital.

Both AI and human intelligence are required. Simply replacing public servants with AI will not always improve outcomes. And yet neither will giving civil servants complete freedom to integrate AI into their own ways of working, with limited guidance or oversight.

Instead of these more extreme approaches — either pursuing wholesale automation or a hands-off approach to adopting AI at personal discretion — attendees believed a balanced approach was essential. Government should develop AI tools which let humans and AI do what they are each best at, freeing public servants' time to focus on core human decision-making tasks.

Government can do this best by embedding these AI tools in cross-functional teams informed by considered design principles. Having effective design principles is key. Recent research has found that human-AI teams can sometimes perform worse than either humans or AI operating alone — the ways human users work with AI is a crucial factor.[8]

A whole system approach

Attendees insisted that, when considering where to use AI, consideration must also be given to the broader public service delivery system. AI may reduce the time it takes to perform and/or increase the quality of a certain service. For example, it could increase the speed and accuracy of chest X-ray analysis. However, if there are blocks at other stages of the delivery system — there may be limited CT scan capacity, delays in communicating with patients, or severely constrained treatment capacity — then overall outcomes will fail to improve.

Government needs to, as one attendee put it, "keep a constant eye on the system as a whole" and consider where AI can make the most difference to overall performance.

[8] Michelle Vaccaro, Abdullah Almaatouq, and Thomas Malone, 'What Combinations of Humans and AI Are Useful: A Systematic Review and Meta-Analysis', *Nature Human Behaviour*, 2024.

2.2 Risk tolerance

Failure

No one wants to see their project fail. And yet for AI to succeed some projects must, on occasion, be allowed to fail. If properly evaluated, failed projects can hold valuable lessons for future ones, ultimately helping to improve services.

Attendees identified an aversion to evaluation within Whitehall, caused by a low tolerance for failure. This low tolerance means that projects are inadequately evaluated because officials often wish to avoid admitting they failed. Rather than stigmatising failure, projects should be treated as learning and evidence-gathering opportunities with a shorter policy feedback loop and standardised evaluation processes.

This approach is essential if AI is to reach its potential in government. By way of example, several attendees discussed their own experiences of when AI projects did not work out the way they had been intended, what they have learned from those experiences, and called for them to be discussed with the same level of openness that successful case studies are.

Risk

Linked to this aversion to evaluation and failure is an aversion to taking risks. Attendees expressed frustration with what they perceived to be an unduly risk-averse attitude within government, particularly visible in the realm of data processing.

When seeking to deploy AI and process data in an automated way it is vital that civil servants comply with all legal requirements. However attendees suggested that legal ambiguities, particularly around GDPR, mean that "people are terrified of processing data in new ways" and err on the side of caution. Not only does this slow the spread of AI but in many areas "doing nothing is a greater risk". One attendee at a recent *Reform* roundtable noted how in children's social services concerns about the risks of using AI to process documents are given greater attention than the risks to children of services continuing to under-perform when caseworkers manually process documents.

Government cannot be afraid of change. The risks of processing data in new ways must be considered, but so too must the risks of continuing to hold data hostage in operational siloes, unable to be used for machine learning because of risks which are already apparent in how frontline public servants use the same data every day.

Idea 1: Require all departments to report on their plans to adopt AI as part of each department's Outcome Delivery Plan, with priority areas and resources being put into delivering the transformation

Idea 2: Create a cross-government competition to identify the best new ideas for adopting AI in Whitehall, building on work already done by 10 Downing Street's Data Science team (10DS).

Idea 3: Introduce a standardised end-of-pilot evaluation process which determines the lessons learned and whether the pilot should advance to wider rollout.

3. The skills to succeed

As important as shifting the culture around AI is ensuring that the civil service is staffed by people who are knowledgeable about AI.

3.1 Skills

External hires

Attendees identified pay as a significant barrier to attracting high-quality external talent into AI roles within Whitehall.

A pay gap between the public and the private sectors exists in most areas, but this gap is particularly pronounced in AI roles. Last year OpenAI researchers were reportedly hired on more than $800,000 a year[9] and "a top quartile engineer with five years of experience in Silicon Valley can get four or five million dollars a year in salary".[10] Whilst comparable UK private sector roles are less than this, the gap to the public sector is still significant enough to make like-for-like competition to recruit in-house AI expertise very challenging.

It is unrealistic to expect the public sector to match these salaries. But there is only so far you can, in the words of Sir John Kingman, "stretch the elastic".[11] If salaries on offer are entirely incomparable then the civil service will be unable to access the skills required to maximise its benefit from AI.

Internal capability

As important as bringing in external talent is developing the civil service's internal talent. Attendees discussed measures the civil service has already taken — for example the *Civil Service People Plan* emphasises digital and data skills training — but they warned that this training must involve non-specialist civil servants within Whitehall and on the frontline of delivery.[12] Within government a higher level of basic digital and technological literacy about AI would make it easier for opportunities and risks to be identified. A higher base awareness of AI in frontline workforces would make it easier for AI to be implemented in those areas with support from staff. Many of the skills needed are not unique to AI. Instead, they involve more analytical and economic capability to both evaluate the outcome changes from pilot projects and to assess the benefits for future business cases.

One roundtable attendee believed that digital training for non-specialist civil servants — particularly those on the frontline — has long been neglected, resulting in a situation where many staff lack the skills and/or the confidence necessary to fully capitalise on AI's benefits. Without confidence that AI can fully navigate the complexity of practice-based roles, public servants will not trust it to automate parts of their workflow. Correcting this must be an urgent priority.

[9] Jo Constanz, 'OpenAI Engineers Earning $800,000 a Year Turn Rare Skillset Into Leverage', *Yahoo Finance*, 22 November 2023.
[10] Civic Future, 'Should the State Play More of a Role in the Economy?', YouTube, 12 July 2024.
[11] Institute for Government, *Why Is Civil Service Reform so Hard? Sir John Kingman in Conversation with Bronwen Maddox*, 2020.
[12] Government People Group, *Civil Service People Plan 2024-2027*, 2024.

3.2 Leadership

Leading

Attendees felt that the Senior Civil Service (SCS) are increasingly enthusiastic about AI and want to see it implemented. And yet, despite this enthusiasm, attendees noted that the SCS still displayed a tendency to maintain a distance and delegate AI decisions to specialist experts. To a degree this is sensible, as deep specialist knowledge will be essential to successful delivery.

But leaders across government, whether they are technical experts or not, need to set the direction of the policy advice they give and the services they oversee. In the same way that civil service performance is not treated as just an issue for HR departments to resolve, AI should not be treated as just an issue for IT departments.

As AI becomes part of more and more government activity, leaders cannot delegate responsibility entirely. Instead, they need to have the functional skills to understand how it is used and set the long-term direction for this: spotting opportunities for cross-departmental collaboration, managing trade-offs, holding others to account for their proactiveness towards AI adoption, and propelling the cultural transformation around AI.

While the SCS' current enthusiasm for AI is positive, it is insufficient without the skills to lead adoption through their day-to-day jobs.

3.3 Frontline users

Benefits

A consistent theme across the roundtables was the importance of listening to those on the frontline of service delivery. They have the best knowledge about operational challenges, can have some of the best suggestions for where AI could improve their work, and are essential for feedback on the effectiveness of tools.

Attendees provided numerous examples of good engagement with frontline staff, which had given them useful feedback and enabled smoother integration of technology into their workflows. One attendee provided the example of social workers using a form of AI in multiple different, unexpected ways from how policymakers had originally envisioned. Another attendee provided the example of Jobcentre Plus employees threatening to quit if the prototype AI tool they had been provided was even temporarily withdrawn. In both these scenarios engaging the frontline provided access to information which improved the performance of the overall service.

How to engage users

It is not enough to simply present tools to public servants who need to use AI, and six months later ask for feedback. Instead, the users must be brought in from the very start of the policymaking process. The first question must be, as one attendee suggested, "what is the frontline already doing that we can build on?" Adopting this user-first approach and targeting AI pilots at areas recommended by those delivering the service would ensure that the most promising areas are identified and addressed first.

Nevertheless, while technologists should listen to users, they should avoid becoming captured by their voices alone. Attendees highlighted how some parts of the public sector had been culturally opposed to previous transformation initiatives. Furthermore, given the dramatic opportunities presented by AI, many may struggle to imagine a new system. As Henry Ford is alleged to have said, "if I had asked people what they wanted, they would have said faster horses".[13] Government should not be afraid of piloting new approaches in the face of incredulity if it believes that, all evidence considered, AI could still improve outcomes.

3.4 Trust and ethics

Building trust

Attendees acknowledged the importance of ensuring that civil servants trust AI whilst also recognising when not to trust AI. One commented that "if [leadership] oversells AI to staff, and it fails on them, then they won't trust it again".

Celebrating AI's successes, and showing how it could benefit civil servants, was a common theme from attendees. If civil servants can see evidence that AI could benefit them they will be more likely to "take the plunge" and use it. This approach should ensure that any qualms are proactively addressed and new pilots are tailored to civil servants' needs, further cementing trust and increasing the likelihood that AI is successfully embedded in their everyday work.

Ethical concerns

There are legitimate ethical concerns about using AI, particularly in some key areas. Machine learning algorithms can, based on their training data and modelling approach, reflect human biases which can potentially lead to discrimination.

There are certain areas where the public are less likely to welcome AI. One attendee discussed their experience of talking to parents about AI in education. Parents were happy for AI to be used when generating lesson content, but less comfortable with it being used to write school reports. Forcing AI adoption in the face of resistance from the public risks backlashes.

And within the public sector, the impact of AI will not be felt equally. Some kinds of roles involve many repeatable tasks which could be easily automated. Other kinds of roles are more flexible, based on professional judgement and situational awareness which cannot be easily replicated by AI.

Government must consider the risks of discrimination, backlashes and job displacements when deploying AI, or face the prospect of making investments which ultimately waste public money and could deliver poor outcomes for service users.

Idea 4: Increase pay allowances when recruiting into AI-related and AI-adjacent positions.

Idea 5: Produce a standardised AI Essentials training course for Grade 6 and 7 civil servants in non-specialist roles covering basic skills and ethical considerations of AI.

[13] Patrick Vlaskovits, 'Henry Ford, Innovation and That "Faster Horse" Quote', *Harvard Business Review*, 29 August 2011.

4. Streamlining processes

For innovations like AI to succeed they need to navigate a complex series of governance approvals. The threshold for evidence and feedback should be high. But attendees felt that current processes were too complex and believed that much simpler processes could maintain similarly high standards. One believed that "simplifying processes is one of the most important things [government] can do" to accelerate the deployment of AI.

Currently, complex guidance on governance, funding and procurement prevent AI projects from being developed quickly, and other priorities overtake the fledgling AI projects which are then accidentally deprioritised.

4.1 Governing

Simplification

Good guidance and governance are necessary to ensure that AI is safe and effective. But the amount of bureaucracy involved has, in the words of one attendee, "accumulated to an embarrassing degree". Guidance is now simultaneously too extensive — there are 29 separate guidance documents which civil servants wanting to adopt digital technology in health and social care must consider — and too little.[14] *The Generative AI Framework for HM Government* is long but relatively vague on how to carry out evaluation of AI projects, or at what stage of the development process different policies should apply.[15]

This presents a barrier to AI adoption, where the complexity of internal governance is likely to dissuade civil servants from experimenting before they even begin. Simplifying this guidance would allow civil servants to develop AI informed and guided — but not unnecessarily constrained — by regulation.

Enabling centre

Even with simplified guidance the responsibility will always be on civil servants to ensure that they are developing software safely and ethically. Attendees feared that this could still pose a barrier to AI adoption without sufficient expert support to help departments navigate the most challenging policy issues surrounding AI, a fast-evolving technology.

The new digital centre of government in DSIT should be an active and enabling body for AI adoption across the civil service. Attendees recommended that this central body needs to proactively reach out to technical teams working on AI projects to support them in complying with the guidance and seek out recommendations for further simplification.

[14] NHS AI and Digital Regulations Service for Health and Social Care, 'All Adopters' Guidance', Web Page, n.d.
[15] HM Government and Central Digital and Data Office, *Generative AI Framework for HM Government*, 2024.

4.2 Funding

Turning pilots into mature products

If AI's benefits are to be realised government must be braver, accepting risk and the occasional failure as par for the course. And attendees were positive about the government's early experiments, funding pilot projects to test AI in different areas.

Nevertheless, a balance must be struck. Funding lots of pilots is a waste of public money if the lessons from them are forgotten. Successful projects need to be built on further to develop them into full enterprise-level tools, and at the same time projects which show limited promise should be scrapped. One potential cause of the current situation are the incentives within government. One official interviewed for a recent *Reform* paper discussed how leaders often feel the pressure to "pilot AI in [their] team because it makes [them] cool, not because it'll do anything".[16] Another potential cause is the lack of funding. Pilots are cheap; scaling up a pilot is expensive.

Attendees wanted the civil service to develop a more coherent "theory of change" focused on adopting mature AI products over continually piloting different experimental solutions. Funding for failed, duplicate or unviable pilots should be diverted towards those with a realistic chance of realising a return on investment. The "Scan → Pilot → Scale" approach outlined in the *AI Opportunities Action Plan* is a positive step in this direction.[17]

Starting small

Attendees cautioned against rushing towards deploying AI everywhere. Given the challenging fiscal situation the government finds itself in it cannot afford to overspend on expensive AI solutions in areas where simpler and cheaper digital solutions are available. Of course there are big opportunities to use AI in complex and challenging areas. But it should first be established whether simpler technology would suffice instead of AI, which is not essential in every area.

4.3 Partnering with the private sector

Copying best practice

Private sector attendees contrasted their own approach towards accessing software — characterised by flexibility and the use of off-the-shelf tools ahead of expensive custom-made products — with Whitehall's approach.

Central government remains attached to the binary choice of either building AI in-house or procuring bespoke products from the private sector. Both approaches can work well,

[16] Joe Hill and Sean Eke, *Getting the Machine Learning: Scaling AI in Public Services* (Reform, 2024).
[17] Clifford, *AI Opportunities Action Plan*.

particularly given the specialist nature of certain parts of the public sector. But both have their weaknesses and challenges.

Often, a hybrid approach of co-producing tools, or procuring off-the-shelf products which are then integrated locally is most suitable. One attendee commented that "for years we have told ourselves that the public sector is special, and in reality we aren't always. Many off-the-shelf products would work just as well for us as they do for private companies". In the right circumstances, government needs to be more willing to consider utilising off-the-shelf or open-source products as a faster and cheaper way to access high-quality software products.

Genuine expertise

The private sector has a crucial role to play in supporting government with technical expertise. But several attendees pointed out the importance of accessing the right technical expertise for the challenges. Currently, 45 per cent of government digital spend is with large IT companies.[18] Working with such companies has benefitted public services in many areas. But the procurement systems are set up in such a lengthy, costly and bureaucratic manner that SMEs (many of whom have much more specialist skillsets) are often closed out. To ensure that the civil service is accessing the best breadth of expertise possible, procurement must be much more flexible and open to competition from SMEs.

Idea 6: Set aside funding for scaling up successful AI pilots rather than investing all funding up-front in separate projects. Initial funding for pilots should be conditional on the existence of an evaluation plan, clear success measures and requirements for scaling-up.

Idea 7: Streamline the procurement process to make it easier for government to access off-the-shelf products and for SMEs to offer their services to government.

[18] Hill and Eke, *Getting the Machine Learning: Scaling AI in Public Services*.

5. Fixing the plumbing

Underpinning AI adoption are the perennial problems of legacy systems, data quality and data access. These have held back earlier waves of digital transformation and, without a realistic strategy for fixing them, AI's effectiveness within government will always be constrained by artificial limits.

5.1 Iterative improvements

Importance

The continued existence of legacy systems, data quality and data access issues significantly impede the widespread rollout of AI tools. AI models cannot be adequately trained without access to high-quality data. And it can be costly and time-consuming to clean and reformat data for training, or to rebuild legacy systems to allow AI models to operate on 'live' data. One official described implementing AI without addressing these issues as "hard, slow and expensive, because we have a skip behind us with some very old stuff in it which we are dragging along".

Departments have historically struggled to obtain funding to remediate legacy IT, despite the importance of doing so and the cost of continuing to run it. Some attendees suggested that this is because these areas are not at the cutting edge of new technology. Civil service leaders and ministers have limited incentives to fund updating legacy IT instead of more innovative projects which are more exciting and eye-catching to announce.

Raising the floor and the ceiling

While these issues need addressing, attendees advised against adopting a linear approach which postpones all AI development until after all legacy technology systems had been remediated.

Fixing the backlog of legacy technology is incredibly challenging: over the past 25 years there have been 11 different strategies focused on the underlying issues.[19] The government cannot afford to simply hope that the twelfth strategy will be automatically successful. Indeed even if it was, the time it would take to remediate all issues means that government would likely "miss the boat" on many of the benefits from AI in the short-term.

Furthermore, digital technology improves iteratively, not in rare 'big bang' moments. As public services and the technology needed to deliver them continue to advance, different pieces of software will need fixing and updating. In the words of one attendee, "some legacy technology, however you define it, is always going to exist. And we need to be capable of living with that".

[19] National Audit Office, *Digital Transformation in Government: Addressing the Barriers to Efficiency*, 2023.

Government needs to recognise that addressing legacy infrastructure, data quality and data access issues is an ongoing, iterative process which can and should take place alongside deploying AI. Attendees suggested that government should begin by focusing on updating legacy infrastructure in those issues where it is the most immediate barrier to the deployment of AI.

Idea 8: Treat legacy system updates and data quality and access improvements as iterative, agile processes to be done alongside expanding the use of AI in frontline services.

6. Conclusion

AI has the potential to dramatically reshape how public services are delivered, unlocking significant productivity gains and efficiency savings. The solutions which already exist show what is possible, and roundtable attendees were confident that there are many more groundbreaking opportunities to use AI in public services.

But realising the potential of AI is reliant on achieving a cultural transformation in how Whitehall — and the public sector as a whole — operates. It is insufficient to simply automate a few processes, recruit a new chief AI officer, publish more abstract guidance or pilot a few headline innovations.

The ideas presented in this paper suggest some ways that this cultural transformation could come about. However, whether leaders are willing to ruthlessly and continually prioritise the cultural changes which will underpin real AI innovation will be the ultimate signal of success or failure.

The Government's *AI Opportunities Action Plan* has laid the groundwork for this change.[20] But leaders must sustain this momentum and rework the culture, structure and processes of Whitehall to best lead AI adoption. Close attention to the lessons of past successes and failures is vital if the Whitehall machine is to transform itself to build a better State.

[20] Clifford, *AI Opportunities Action Plan*.

Bibliography

Civic Future. ‘Should the State Play More of a Role in the Economy?’ Web Page, 29 June 2024.

Clifford, Matt. *AI Opportunities Action Plan*. Department for Science, Innovation and Technology, 2024.

Constanz, Jo. ‘OpenAI Engineers Earning $800,000 a Year Turn Rare Skillset Into Leverage’. *Yahoo Finance*, 22 November 2023.

Derby City Council. ‘Next Steps for AI Technology to Protect Services for Derby Residents’. Web Page, 18 December 2023.

Government People Group. *Civil Service People Plan 2024-2027*, 2024.

Hill, Joe, and Sean Eke. *Getting the Machine Learning: Scaling AI in Public Services*. Reform, 2024.

HM Government, and Central Digital and Data Office. *Generative AI Framework for HM Government*, 2024.

Institute for Government. *Why Is Civil Service Reform so Hard? Sir John Kingman in Conversation with Bronwen Maddox*, 2020.

Iosad, Alexander, David Railton, and Tom Westgarth. *Governing in the Age of AI: A New Model to Transform the State*. Tony Blair Institute for Global Change, 2024.

Local Government Association and NHS England. ‘Artificial Intelligence Use Cases’. Web Page, n.d.

Ministry of Defence. *The Defence AI Playbook*, 2024.

National Audit Office. *Digital Transformation in Government: Addressing the Barriers to Efficiency*, 2023.

NHS AI and Digital Regulations Service for Health and Social Care. ‘All Adopters’ Guidance’. Web Page, n.d.

NHS Cambridge University Hospitals. ‘AI Cuts Waiting Times for Cancer Patients in NHS First’. Web Page, 27 June 2023.

Vaccaro, Michelle, Abdullah Almaatouq, and Thomas Malone. ‘What Combinations of Humans and AI Are Useful: A Systematic Review and Meta-Analysis’. *Nature Human Behaviour*, 2024.

Vlaskovits, Patrick. ‘Henry Ford, Innovation and That “Faster Horse” Quote’. *Harvard Business Review*, 29 August 2011.

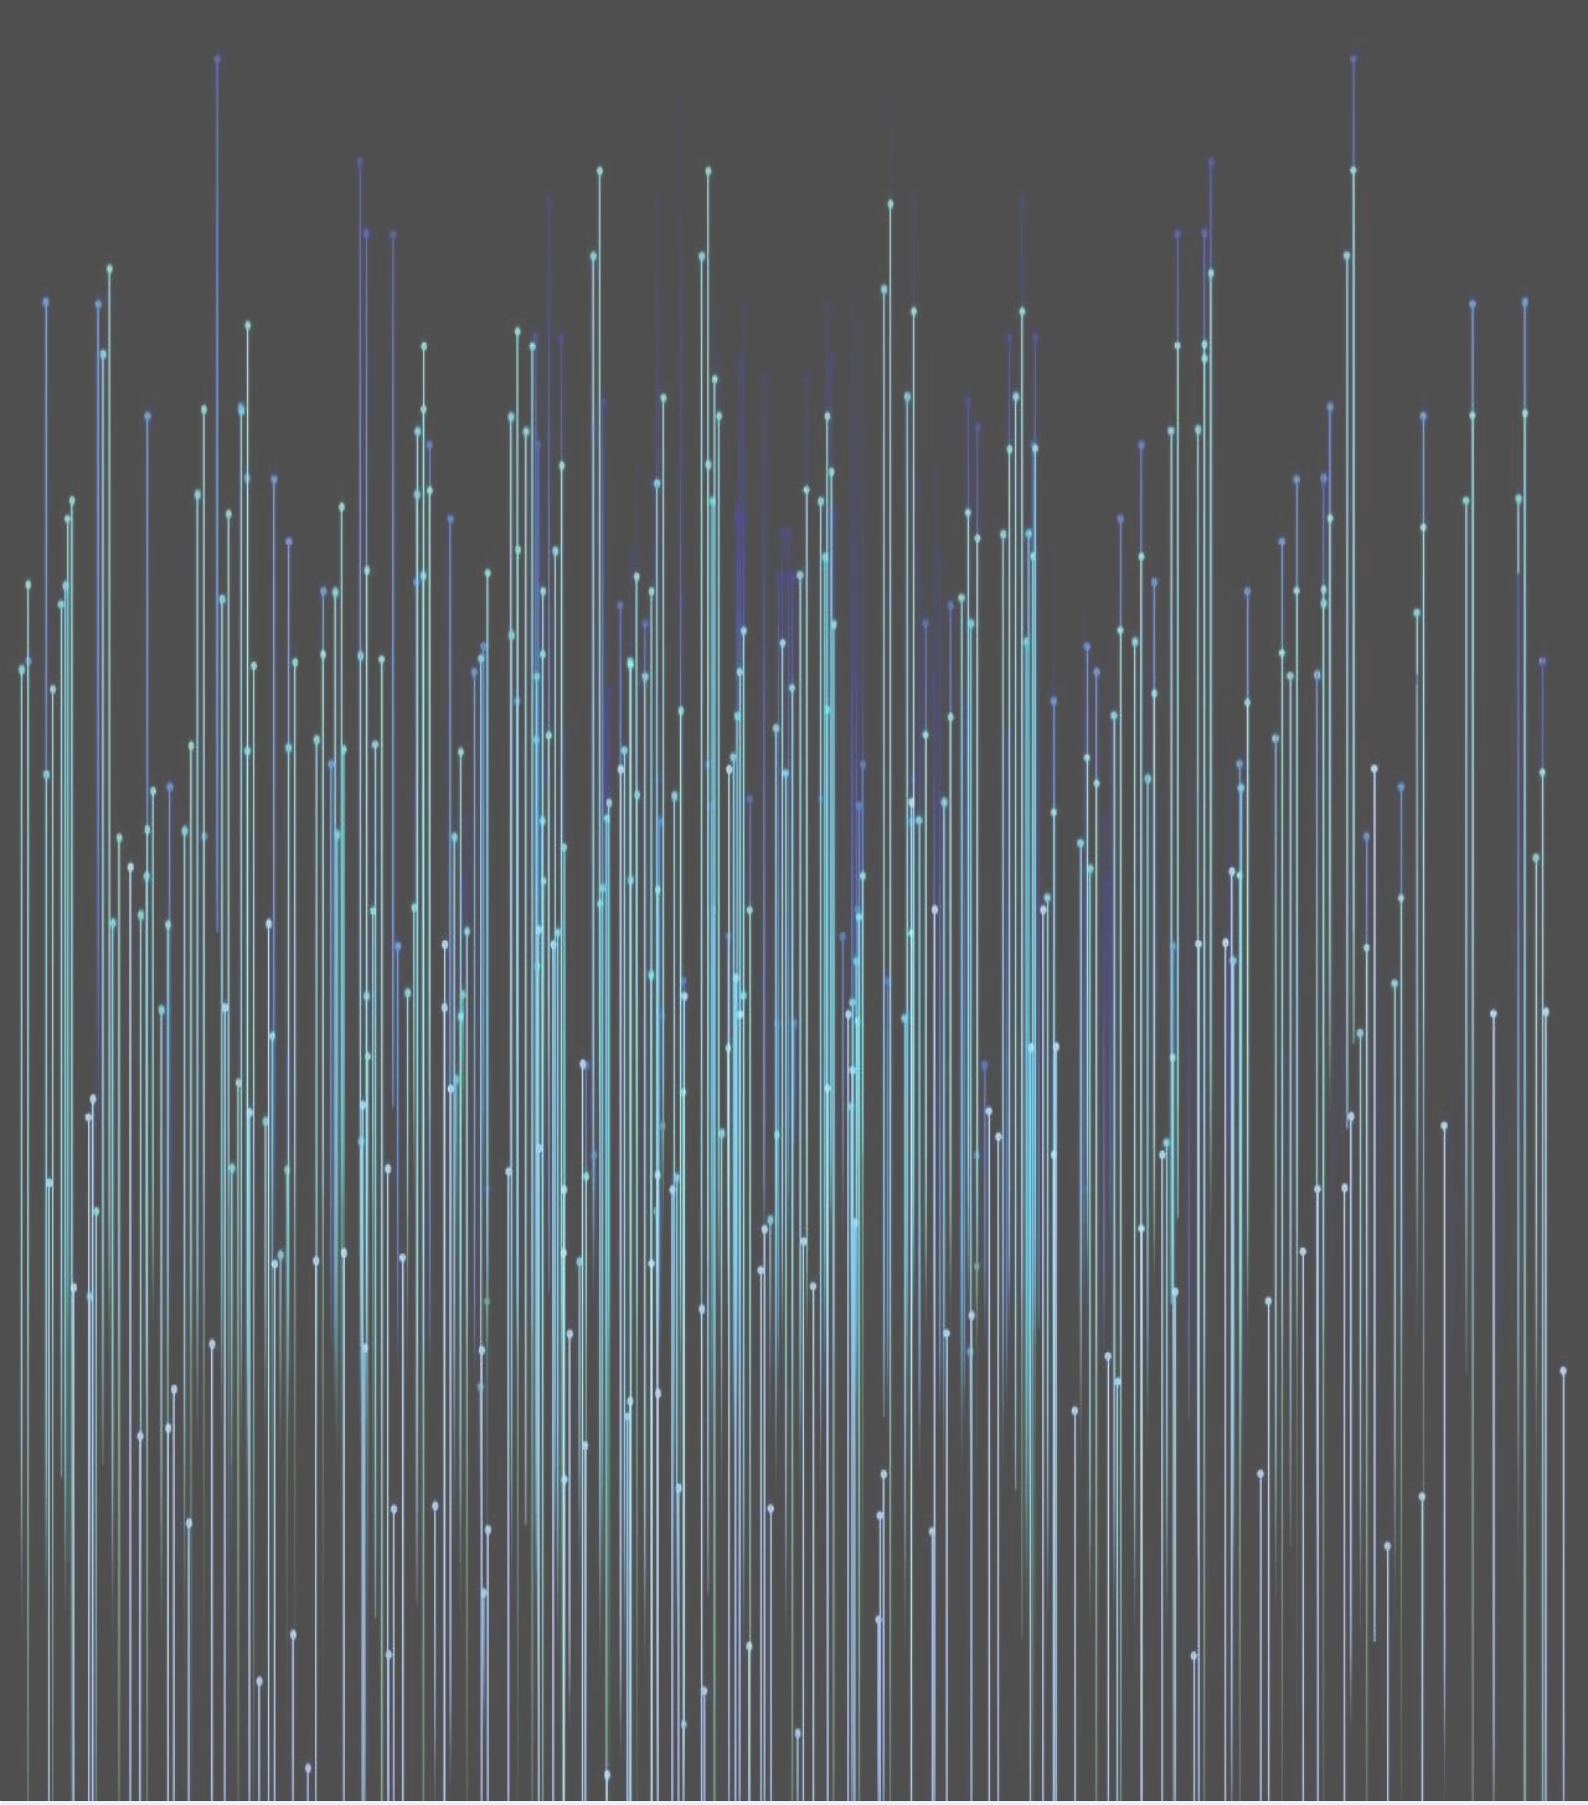

REFORM

978-1-910850-72-5

@reformthinktank

@reformthinktank.bsky.social

www.reform.uk